

# Popularity Feedback Constrains Innovation in Cultural Markets

Lucas Gautheron<sup>1</sup>, Raja Marjeh<sup>2</sup>, Dalton C. Conley<sup>3</sup>, Seth Frey<sup>4</sup>, Hannah Rubin<sup>5</sup>, Mike D. Schneider<sup>5</sup>, Ofer Tchernichovski<sup>6</sup> & Nori Jacoby<sup>7,8</sup>

<sup>1</sup>Evolution, Science and Society, University of Missouri, Columbia, Missouri, United States  
(lucas.gautheron@gmail.com)

<sup>2</sup>Department of Psychology, Princeton University, Princeton, New Jersey, United States

<sup>3</sup>Department of Sociology, Princeton University, Princeton, New Jersey, United States

<sup>4</sup>Department of Communication, UC Davis, Davis, California, United States

<sup>5</sup>Department of Philosophy, University of Missouri, Columbia, Missouri, United States

<sup>6</sup>Department of Psychology, Hunter College, New York, New York, United States

<sup>7</sup>Cornell Computational Cognition Lab, Cornell University, Ithaca, New York, United States

<sup>8</sup>Max Planck Institute for Empirical Aesthetics, Frankfurt, Germany

## Abstract

Real-world creative processes ranging from art to science rely on social feedback-loops between selection and creation. Yet, the effects of popularity feedback on collective creativity remain poorly understood. We investigate how popularity ratings influence cultural dynamics in a large-scale online experiment where participants ( $N = 1\,008$ ) iteratively *select* images from evolving markets and *produce* their own modifications. Results show that exposing the popularity of images reduces cultural diversity and slows innovation, delaying aesthetic improvements. Popularity feedback is associated with changes to both selection and creative stages. During selection, popularity information triggers cumulative advantage, with participants preferentially building upon popular images, reducing diversity. During creation, participants make less disruptive changes, and are more likely to expand existing visual patterns. Feedback loops in cultural markets thus not only shape selection, but also, directly or indirectly, the form and direction of cultural innovation.

**Keywords:** cultural evolution; collective creativity; feedback loops; cumulative advantage; cultural runaway

## Introduction

In many domains of human culture such as science, arts, or social media, producers of cultural goods can observe and compare the popularity of existing cultural variants. These popularity signals can trigger a form of *popularity feedback* by influencing later cultural innovations. In science, for instance, popularity information is widely available in the form of citation counts, which are known to exhibit cumulative advantage as scientists preferably cite highly popular papers (Barabási & Albert, 1999; Wang et al., 2013). On the producer side, disproportionate attraction to popular ideas may diminish intellectual diversity, distort assessments of the long-term potential of certain research directions, and alter the space of possibilities that scientists are willing to explore (Kummerfeld & Zollman, 2016). Alternatively, popularity feedback may enhance social learning by providing reliable fitness signals (Miu et al., 2018). Additionally, popularity feedback may alter innovation strategies in ways unappreciated by blind variation models (Campbell, 1960), altering the shape of cultural products. For instance, in artistic domains, we can hypothesize that access to popularity feedback encourages producers

to amplify or exaggerate arbitrary features present in successful cultural goods. This could create a “cultural runaway” (Boyd & Richerson, 1988) analogue to biological runaway selection, an evolutionary process responsible for exaggerated ornaments across several species and driven by a positive feedback between a trait and preference for this trait (Lande, 1981; Lehtonen & Kokko, 2012; Prum, 2012). Similarly, digital platforms such as social media may be influenced by popularity feedback via shares or likes. While no evidence of such popularity biases was found in Twitter (Carrignon et al., 2019; Youngblood et al., 2023), these could arise in platforms with higher stakes for producers.

Despite its potential significance, few works have addressed the evolutionary consequences of popularity feedback. Prior research has often neglected the incidence of feedback loops on the form and direction of cultural innovations, owing to a stronger focus on selection in cultural evolutionary literature (André et al., 2023). Better studied feedback mechanisms include frequency-dependent selection, when the adoption rate of cultural variants depends on their incidence among a population, due to e.g. conformity biases (Acerbi & Alexander Bentley, 2014; Newberry & Plotkin, 2022), or varying utility, as in the case of “network goods” (Björkegren, 2018). A related phenomenon is cumulative advantage, when popular cultural goods that are already popular attract more attention, independently from their intrinsic worth. For instance, Salganik et al. (2006) have shown that access to popularity information in artificial markets of songs ultimately led to more unequal and unpredictable market shares. By contrast, the effects of such feedback loops in markets involving both consumers *and* producers of cultural goods has been far less studied experimentally, with the exception of Baliotti et al. (2016), who demonstrated that competitive incentives influence peer-review feedback processes involving both creators and reviewers. Yet, this study did not consider social feedback via popularity information, and it did not directly examine how social information affects the shape of cultural outputs.

We address these gaps in a large-scale online experiment in which participants ( $N = 1\,008$ ) iteratively observe a cultural market of images and populate it with their own drawing after

having selected an image to improve. Our experiment thus implements both selection and creation steps, therefore capturing the two major forces of cultural evolution (Mesoudi, 2021). We find that access to popularity information (about how many times items were copied) leads to less diverse cultural markets; slower innovation; and initially less aesthetic and creative images, although this trend may be reversed in the long-run, possibly due to social feedback protecting high-quality innovations. These effects are associated with changes to both selection and creation steps. During selection, popularity information triggers a process of cumulative advantage, whereby participants preferentially choose to build upon popular ideas, which undermines diversity and amplifies popularity inequalities. Finally, individuals make less disruptive changes under popularity feedback. We discuss what can possibly lead to these changes in the emergent cultural dynamics.

## Methods

We develop a large-scale collective experiment in which participants observe collections of images ( $16 \times 16$  grids with black or white pixels), select one image to modify, and contribute their creation to the collection (Fig. 1). The editing interface was similar to other pixel-based drawing experiments in the literature (Hart et al., 2017; Kumar et al., 2022; Tchernichovski et al., 2025). In the treatment group, during the selection step, participants observe the popularity of each item in the market (i.e. how many times the item had been previously selected to be edited). Participants are asked to behave so as to maximize the chance that their own work is selected by subsequent participants. This creates a strategic tension: while participants may choose to capitalize on ideas that have proven successful in the past, doing so implies competing with many similar images for attention, decreasing the chance of being selected. This dilemma thus invites either conformist strategies (building on popular trends, i.e. positive feedback) or anti-conformist approaches (pursuing novel ideas, i.e. negative feedback). The process mirrors dynamics in many creative cultural domains, including science, where researchers must balance engaging with trending topics against differentiating their contributions. After selecting an image, participants are allowed to change no less than 1 and up to 24 pixels. This constraint ensures that participants engage with the selected image and thus increases the strategic importance of the selection step. In addition, the editing restriction dilutes the relative effect of individual contributions.

We consider 128 isolated markets, initialized with a single image each, and evolving as independent chains over 60 generations (7 680 images overall). For every chain, at any generation  $g$ , one participant observes a market of up to  $T = 12$  images produced between generations  $g - 1$  and  $g - T$ . Chains are divided into two conditions: the popularity information condition (PI), in which participants observe the popularity of each image in the market during the selection step (i.e., how many times the image was already selected and edited by previous participants); and the no-popularity information

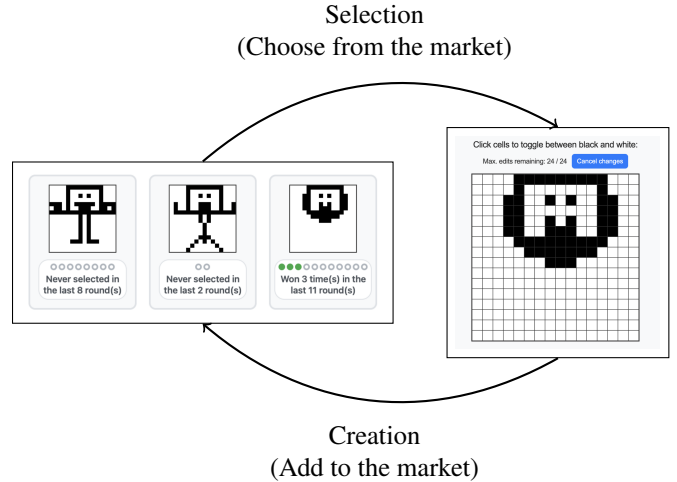


Figure 1: Selection and creation steps. Participants observe a market of 12 images, choose one image, change  $\geq 1$  and  $\leq 24$  pixels, and add their own creation to the market.

condition (NPI), in which this feedback is absent. Every chain from either condition is paired with a chain from the other condition, initialized with the same image (as in Fig. 2). Participants contributed only once for each chain and up to 8 chains overall, all from the same condition (fewer if they dropped out before the end of the experiment).

The experiment was deployed online using PsyNet (Harrison et al., 2020). It includes  $N = 2015$  participants split between the cultural market experiment itself ( $N = 1008$ ) and additional image rating or labeling experiments ( $N = 1007$ ). Participants were recruited in the United States via Prolific, provided informed consent under a Cornell University-approved protocol (IRB0148995) or a Max Planck Society-approved protocol (2021-4/2024-36), and were compensated at a rate of approximately \$12 per hour.

## Results

### Access to popularity information undermines cultural diversity and constrains innovation

We examine the consequences of popularity information for cultural diversity and innovation. First, at each generation of every chain, we measure cultural diversity as the average pairwise distance between images in the market, according to three distance measures: the phylogenetic distance (i.e., the degree of separation between images in their evolutionary tree); the pixel-wise (or hamming) distance, measuring the fraction of pixels that differ between two images; and the semantic distance between images (i.e., the cosine distance between high-dimensional image-embeddings computed by a vision-transformer model; Li et al. 2022). This reveals a statistically significant reduction in diversity across the three measures when social feedback is provided (Fig. 3).

Popularity information may affect not only the state of cultural markets at any point in time, but also the rate of cultural

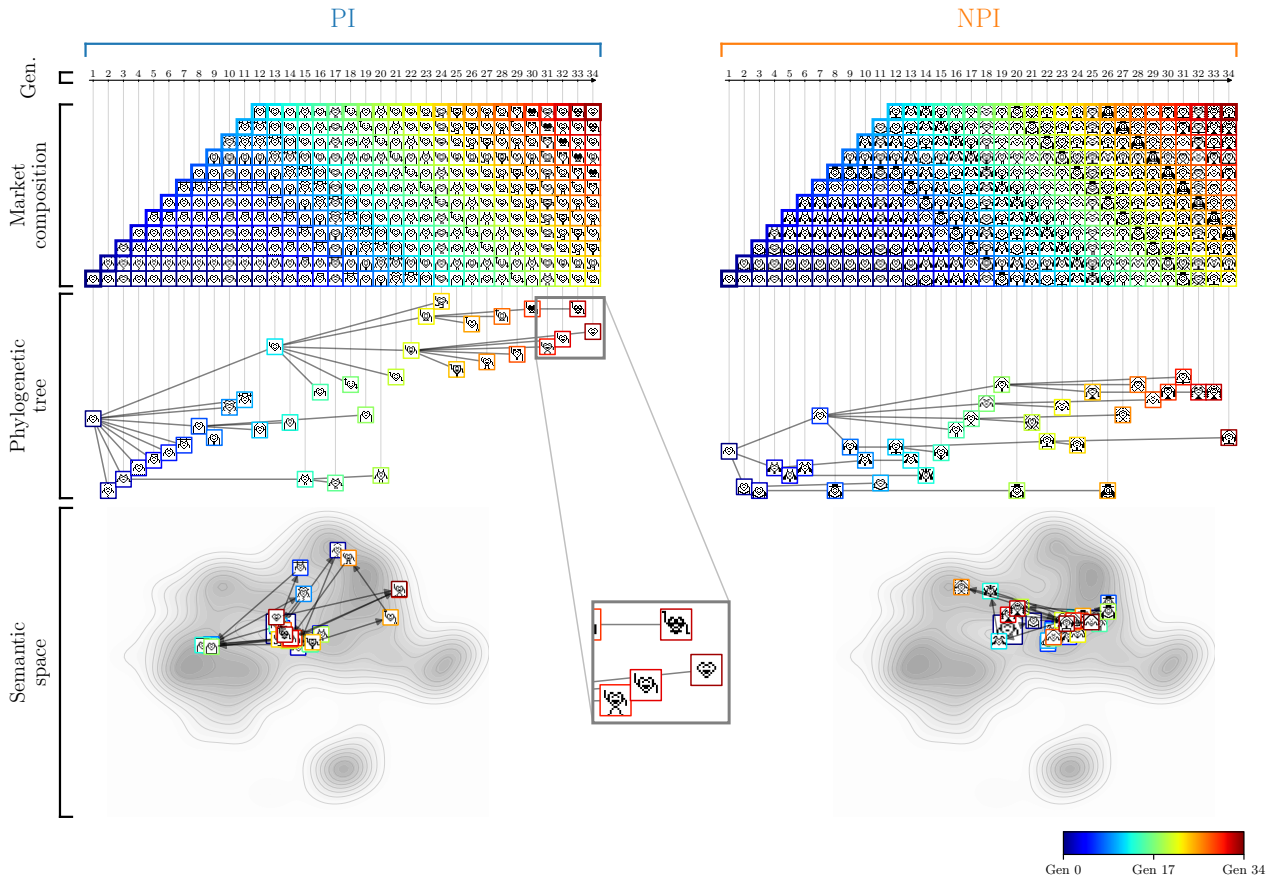


Figure 2: Evolution of two chains starting from the same image, over the first 35 generations, under conditions with popularity information (PI, to the left) and without (NPI, to the right). **Market composition:** at each generation (*gen.*), participants observe a market containing the 12 latest creations in the chain (each column represents the contents of the market at a given generation). **Phylogenetic tree:** ancestry relationships between images can be represented in the form of a tree, from the first image (to the left,  $t = 1$ ) to the last (on the right,  $t = 34$ ). **Semantic space:** images are assigned positions in a high-dimensional space using a machine learning vision model (Li et al., 2022), revealing the motion of each chain in their broader cultural landscape.




Figure 3: Average market diversity in the conditions with popularity information (PI) and without (NPI). Diversity is measured as the average distance between images in a market at any generation. Error bars indicate standard errors ( $\pm 1SE$ ). Statistical significance is evaluated via a non-parametric permutation test pairing chains with identical initial images.




Figure 4: Average chain autocorrelation in semantic space  $\bar{\rho}(\tau)$  under the conditions with popularity information (PI) and without (NPI) and for different time-lags  $\tau$ .

innovation across time. In order to estimate the effect of popularity information on temporal exploration dynamics, we evaluate the average chain autocorrelation  $\bar{\rho}(\tau)$  as follows:

$$\bar{\rho}(\tau) = \frac{\sum_{c,t} \mathbf{x}'_{c,t} \cdot \mathbf{x}'_{c,t+\tau}}{\sum_{c,t} \mathbf{x}'_{c,t} \cdot \mathbf{x}'_{c,t}} \quad (1)$$

where  $\mathbf{x}'_{c,t} = \mathbf{x}_{c,t} - \boldsymbol{\mu}$  is the position (in semantic space, cf. Fig. 2) of the image from chain  $c$  created at generation  $t$  after subtracting the mean coordinates  $\boldsymbol{\mu} \equiv \frac{1}{C \times T} \sum_{c,t} \mathbf{x}_{c,t}$ , and the dot product  $(\cdot)$  measures the similarity between any image-pair. For a given time lag  $\tau$ , the autocorrelation  $\bar{\rho}(\tau)$  measures the relative similarity between images at time  $t$  and  $t + \tau$  within a chain. High autocorrelation values imply that it takes many generations to visit regions farther away in semantic space, such that it eventually takes longer to explore the underlying cultural landscape. Fig. 4 shows that popularity information increases autocorrelation, which means initial conditions have a longer-lasting impact on the exploration process.

Finally, we hypothesized that feedback loops owing to popularity information may have detrimental effects on the quality of cultural innovation. Quality was measured by collecting appeal, aesthetics, and creativity scores between 1 and 9, for every image, among a new set of 486 human participants (Fig. 5). While differences in ‘‘appeal’’ were small, during most of the experiment (gen. 0–47), images in condition PI are less aesthetic ( $p = 0.002$ ) and creative ( $p < 0.001$ ). Towards the end of the experiment (gen. 48–59), however, we find this trend to be reduced and/or reversed. While further data is required to confidently ascertain the robustness of this inversion ( $p = 0.14$  for appeal and  $p = 0.14$  also for aesthetics), a potential explanation is proposed at the end of the paper.

In what follows, we show that these effects of social feedback are associated with changes in both selection processes, namely how participants choose ideas to build upon, and creation processes, namely how participants innovate

**Social feedback and selection** We begin by exploring the impact of popularity information on selection. Using individual selection data, and for each condition (PI and NPI), we infer the prevalence of each of four selection strategies: image-driven selection (assuming choices are dictated by a latent intrinsic fitness); cumulative advantage (one of the most popular images is chosen); balancing selection (one of the least popular images is chosen); and uniformly random selection. In image-driven selection, we assume that participants select an image  $k \sim \text{Categorical}(\mathbf{p})$  where  $\mathbf{p} = \text{softmax}(\mathbf{u})$  and  $\mathbf{u}$  is the latent utility of each image on the market. Latent utilities are assumed to balance four criteria  $c$  (‘‘appeal’’, ‘‘editing potential’’, ‘‘originality’’, ‘‘recognizability’’) measured by latent scores  $\mathbf{u}_c$ . These scores are indirectly measured by collecting subjective ratings  $r$  (from an additional new group of 413 participants) and assuming that  $r_{ic} \sim \mathcal{N}(u_{ic}, \sigma_c)$ . We then further assume that  $\mathbf{u} = \sum \beta_c \mathbf{u}_c$ .

We fit this hierarchical model in Stan. This reveals a strong contribution of cumulative advantage in condition PI and its

quasi-absence in condition NPI (Fig. 6). According to the model, balancing selection does not significantly arise; however, preferences for least popular variants could be mistaken as random selection. Overall, these observations are consistent with the reduction in diversity observed under social feedback. Interestingly, in condition PI, cumulative advantage arises at the expense of random selection, rather than image-driven selection (which does not significantly increase in condition NPI). Selection policy mixtures inferred from individual choices can be used to simulate the evolution of phylogenetic trees; this predicts lower phylogenetic under social feedback, consistent with the data. To further assess the robustness of our model, we analyzed the responses of participants in the original experiment, who were asked to self-report their selection strategy, and matched these responses to the mixture of policies inferred from their behavior. Using BERTopic (Grootendorst, 2022), the self-reports were clustered into 7 topics, each of which was then labeled with GPT-5.2. Table 1 shows that Topic 2, which emerges in condition PI and is dominated by reports referring to items’ popularity, is strongly associated with cumulative advantage.

Finally, like Epstein et al. (2021) and Salganik et al. (2006), we compared the inequality in the success of cultural items under the conditions with and without social feedback, in terms of the Gini coefficient (Bendel et al., 1989). We find  $G(PI) = 0.69$  and  $G(NPI) = 0.61$ ; in other words, in line with previous works, social feedback increases success inequalities ( $\Delta = 0.08$ ,  $p < 0.001$ ).

**Social feedback and creation** During creation, under social feedback, participants change fewer pixels to the images they select and make smaller semantic shifts on average (Fig. 7). However, summary statistics such as pixels added and removed say little about the underlying strategy. To get further insight into how popularity information affects the creation process itself, we classified edits into five categories (Fig. 8): ‘‘disruption’’ (when the new image represents something fundamentally different); ‘‘addition’’ (when visual features are added without changing the meaning of the image); ‘‘pattern-growth’’ (when existing features are expanded); ‘‘removal’’ (when existing features are removed, without changing the meaning of the image); and ‘‘refinement’’ (any other minor type of change). A new group of 372 humans were recruited online to classify parent-child image pairs into either of these categories, based on the majority-choice among three annotations per pair. We then calculated the prevalence of each strategy across conditions (Fig. 9). We find that in the condition with popularity information, participants are indeed less likely to disrupt the original idea (OR=0.84, CI<sub>95%</sub> = [0.73, 0.96]). Instead, participants are more likely to grow and expand existing visual features (OR=1.18, CI<sub>95%</sub> = [1.02, 1.36]), in line with the runaway hypothesis. Evidence also suggests that social feedback makes the removal of visual features less likely (OR=0.82, CI<sub>95%</sub> = [0.66, 1.00]).




Figure 5: Quality of images throughout generations and across conditions. Error bars indicate standard errors ( $\pm 1SE$ ) on the chain averages (effective sample-size  $n = 64$ ).




Figure 6: Left & middle plots: effect of popularity information on selection. Right plot: phylogenetic diversity predicted by the best-fit mixture of selection strategies, for each condition. Error bars represent 68% credible intervals (left & middle plots) and standard errors (right plot).

Table 1: Self-reported image selection strategies (clustered into 7 topics labeled with GPT-5.2), on a subset of the data (50%). Columns indicate the number of participants in each cluster ( $N$ ), the average prevalence of selection policies among these participants, and the distribution of these participants in each condition.

Self-reported strategy cluster	$N$	Image-driven	Cum. adv.	Bal. sel.	Random.	PI	NPI
0_Appealing and editable	140	59%	13%	2%	26%	47%	53%
1_Editability improvement potential	127	60%	14%	0%	26%	39%	61%
2_Popularity and past wins	51	37%	43%	1%	19%	86%	14%
3_No Deliberate Strategy	42	56%	9%	2%	34%	43%	57%
4_Popularity and editability	48	58%	25%	2%	16%	50%	50%
5_Aesthetic Appeal and Editability	58	58%	18%	1%	23%	50%	50%
6_Gut appeal and potential	20	62%	8%	5%	26%	55%	45%




Figure 7: Average edits magnitude in each condition.




Figure 8: Illustrative examples for each of the five editing strategies.




Figure 9: Prevalence of editing strategies across conditions. Error bars indicate 68% CIs, under uniform priors.  $p$  is the probability that the posterior includes the null (OR=1).

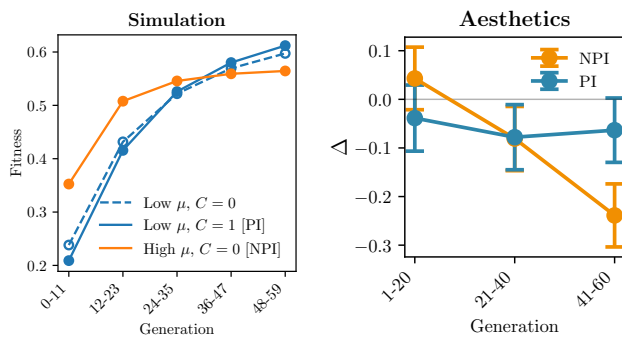


Figure 10: Left: evolution of images' fitness in an idealized selection/mutation model, under low versus high mutation rates, and in the occurrence of cumulative advantage. Social feedback (PI) is captured by low mutation rates and cumulative advantage (low  $\mu$ ,  $C = 1$ ), versus high mutation rates and no cumulative advantage for NPI (high  $\mu$ ,  $C = 0$ ). Right: aesthetic improvements across generations and conditions in the data.  $\Delta$  is the difference between the aesthetic scores of edited images and their parents.

**Aesthetics and popularity feedback** Finally, we discuss why images might possibly be initially less but ultimately more aesthetic under social feedback (explaining the inversion seemingly occurring in the last iterations, cf. Fig. 5). We consider a model where innovations are represented by bitstrings  $(b_i) \in \{0, 1\}^N$  with fitness  $f = \sum_i b_i$  (MacKay, 2003, p. 269). At each round, with 50% probability, individuals select the image with highest fitness in the market (consistent with the measured rate of image-driven selection, Fig. 6), or otherwise choose at random. Then, they transform the image by mutating a few pixels. For higher mutation rates, as in condition NPI, the model predicts that images improve faster in the first generations (Fig. 10). This is because when images are far from optimal, most mutations are beneficial. However, high mutation rates fare worse at later generations, when most changes are detrimental. The data indeed shows initially better but gradually worse variations in aesthetics in condition NPI, consistent with this model (Fig. 10). Social feedback may thus have a protective effect that initially impedes improvements but ultimately preserves quality. While differences in mutation rates alone may explain an early disadvantage/late advantage under social feedback, cumulative advantage (present in condition PI) reinforces this effect (cf. the dashed versus blue solid lines in Fig. 10).

## Discussion

Our findings demonstrate that positive feedback loops in culture have distinct effects on selection and innovation, with measurable consequences for the diversity and quality of the cultural goods. Selection is distinctively influenced, with significant occurrence of cumulative advantage under social feedback, which ultimately reduces the diversity of cultural markets. By studying an open-ended creative task, we were additionally able to observe the diversity of innovation strategies and their alteration under the influence of social feedback: innovations were more conservative/meaning-preserving, and more likely to expand existing features, consistent with a cultural runaway process.

Future work, however, should address several limitations. First, control experiments should determine whether changes in innovation strategies partly follow from changes in the distribution of images being selected, which could be achieved by ablation studies similar to those from Marjeh et al. (2025). Second, participants did not observe the success of their own images; it is conceivable that they would adjust their strategies otherwise. Third, the constraint on editing, which intends to limit individual contributions, may not suffice to capture cumulative cultural evolution (Miton & Charbonneau, 2018). Finally, in spatially extended systems, partial coupling between neighboring markets may alleviate or amplify some effects of social feedback, which could be addressed in a network-based experiment (Shiiku et al., 2025). More broadly, our work demonstrates how innovations in large-scale online simulations of cultural evolution can inform our understanding of emergent features in cultural markets.

## Acknowledgments

This work was supported by the NSF grant “Collaborative Research: Designing smart environments to augment collective learning & creativity” (award 171418).

## References

- Acerbi, A., & Alexander Bentley, R. (2014). Biases in cultural transmission shape the turnover of popular traits. *Evolution and Human Behavior*, 35(3), 228–236. <https://doi.org/10.1016/j.evolhumbehav.2014.02.003>
- André, J.-B., Baumard, N., & Boyer, P. (2023). Cultural evolution from the producers’ standpoint. *Evolutionary Human Sciences*, 5. <https://doi.org/10.1017/ehs.2023.20>
- Baliotti, S., Goldstone, R. L., & Helbing, D. (2016). Peer review and competition in the art exhibition game. *Proceedings of the National Academy of Sciences*, 113(30), 8414–8419. <https://doi.org/10.1073/pnas.1603723113>
- Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *science*, 286(5439), 509–512.
- Bendel, R., Higgins, S., Teberg, J., & Pyke, D. (1989). Comparison of skewness coefficient, coefficient of variation, and gini coefficient as inequality measures within populations. *Oecologia*, 78(3), 394–400.
- Björkegren, D. (2018). The adoption of network goods: Evidence from the spread of mobile phones in rwanda. *The Review of Economic Studies*, 86(3), 1033–1060. <https://doi.org/10.1093/restud/rdy024>
- Boyd, R., & Richerson, P. J. (1988). *Culture and the evolutionary process*. University of Chicago press.
- Campbell, D. T. (1960). Blind variation and selective retentions in creative thought as in other knowledge processes. *Psychological Review*, 67(6), 380–400. <https://doi.org/10.1037/h0040373>
- Carrignon, S., Bentley, R. A., & Ruck, D. (2019). Modelling rapid online cultural transmission: Evaluating neutral models on twitter data with approximate bayesian computation. *Palgrave Communications*, 5(1). <https://doi.org/10.1057/s41599-019-0295-9>
- Epstein, Z., Groh, M., Dubey, A., & Pentland, A. (2021). Social influence leads to the formation of diverse local trends. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2), 1–18. <https://doi.org/10.1145/3479553>
- Grootendorst, M. (2022). Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*.
- Harrison, P., Marjeh, R., Adolphi, F., van Rijn, P., Anglada-Tort, M., Tchernichovski, O., Larrouy-Maestri, P., & Jacoby, N. (2020). Gibbs sampling with people. *Advances in neural information processing systems*, 33, 10659–10671.
- Hart, Y., Mayo, A. E., Mayo, R., Rozenkrantz, L., Tandler, A., Alon, U., & Noy, L. (2017). Creative foraging: An experimental paradigm for studying exploration and discovery. *PloS one*, 12(8), e0182133.
- Kumar, S., Correa, C. G., Dasgupta, I., Marjeh, R., Hu, M. Y., Hawkins, R., Cohen, J. D., Narasimhan, K., Griffiths, T., et al. (2022). Using natural language and program abstractions to instill human inductive biases in machines. *Advances in Neural Information Processing Systems*, 35, 167–180.
- Kummerfeld, E., & Zollman, K. J. S. (2016). Conservatism and the scientific state of nature. *The British Journal for the Philosophy of Science*, 67(4), 1057–1076. <https://doi.org/10.1093/bjps/axv013>
- Lande, R. (1981). Models of speciation by sexual selection on polygenic traits. *Proceedings of the National Academy of Sciences*, 78(6), 3721–3725. <https://doi.org/10.1073/pnas.78.6.3721>
- Lehtonen, J., & Kokko, H. (2012). Positive feedback and alternative stable states in inbreeding, cooperation, sex roles and other evolutionary processes. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1586), 211–221. <https://doi.org/10.1098/rstb.2011.0177>
- Li, J., Li, D., Xiong, C., & Hoi, S. (2022). Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. <https://doi.org/10.48550/ARXIV.2201.12086>
- MacCallum, R. M., Mauch, M., Burt, A., & Leroi, A. M. (2012). Evolution of music by public choice. *Proceedings of the National Academy of Sciences*, 109(30), 12081–12086. <https://doi.org/10.1073/pnas.1203182109>
- MacKay, D. J. (2003). *Information theory, inference and learning algorithms*. Cambridge university press.
- Marjeh, R., Anglada-Tort, M., Griffiths, T. L., & Jacoby, N. (2025). Characterizing the interaction of cultural evolution mechanisms in experimental social networks. *arXiv preprint arXiv:2502.12847*.
- Mesoudi, A. (2021). Cultural selection and biased transformation: Two dynamics of cultural evolution. *Philosophical Transactions of the Royal Society B*, 376(1828). <https://doi.org/10.1098/rstb.2020.0053>
- Miton, H., & Charbonneau, M. (2018). Cumulative culture in the laboratory: Methodological and theoretical challenges. *Proceedings of the Royal Society B: Biological Sciences*, 285(1879), 20180677. <https://doi.org/10.1098/rspb.2018.0677>
- Miu, E., Gulley, N., Laland, K. N., & Rendell, L. (2018). Innovation and cumulative culture through tweaks and leaps in online programming contests. *Nature Communications*, 9(1). <https://doi.org/10.1038/s41467-018-04494-0>
- Newberry, M. G., & Plotkin, J. B. (2022). Measuring frequency-dependent selection in culture. *Nature Human Behaviour*, 6(8), 1048–1055. <https://doi.org/10.1038/s41562-022-01342-6>
- Prum, R. O. (2012). Aesthetic evolution by mate choice: Darwin’s really dangerous idea. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1600), 2253–2265.
- Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an arti-

- ficual cultural market. *Science*, 311(5762), 854–856. <https://doi.org/10.1126/science.1121066>
- Shiiku, S., Marjeh, R., Anglada-Tort, M., & Jacoby, N. (2025). The dynamics of collective creativity in human-ai hybrid societies. *arXiv preprint arXiv:2502.17962*.
- Tchernichovski, O., Globerson, E., Harrison, P., & Jacoby, N. (2025). How constraints on editing affects cultural evolution. *arXiv preprint arXiv:2502.16694*.
- Wang, D., Song, C., & Barabási, A.-L. (2013). Quantifying long-term scientific impact. *Science*, 342(6154), 127–132. <https://doi.org/10.1126/science.1237825>
- Youngblood, M., Stubbersfield, J. M., Morin, O., Glassman, R., & Acerbi, A. (2023). Negativity bias in the spread of voter fraud conspiracy theory tweets during the 2020 us election. *Humanities and Social Sciences Communications*, 10(1). <https://doi.org/10.1057/s41599-023-02106-x>